
Chapter 4

Cascades and herds

One million people cannot be wrong.

Each agent observes what others do and takes a zero-one decision in a pre-ordered sequence. In a cascade, all agents herd on a sufficiently strong public belief and there is no learning. In a herd, all agents turn out to take the same decision. A cascade generates a herd but the converse is not true. Cascades are non generic for atomless distributions of beliefs while a herd always takes place, eventually! Since a herd does take place eventually, the probability that it is broken must converge to zero. Hence, there is some learning in a herd (it is not broken), but the learning is very slow. The stylization of that property is the cascade.

Beliefs converge to the truth only if the distribution of private beliefs is unbounded, but the self-defeating principle in social learning implies that the convergence is slow. Since the filter imposed by discrete actions is coarse, the slowdown of social learning is much more significant than in the previous chapter. Applications for welfare properties and pricing policies by a monopoly are discussed.

A tale of two restaurants

Two restaurants face each other on the main street of a charming alsatian village. There is no menu outside. It is 6pm. Both restaurants are empty. A tourist comes down the street, looks at each of the restaurants and goes into one of them. After a while, another tourist shows up, evaluates how many patrons are already inside by looking through the stained glass windows—these are alsatian *winstube*—and chooses one of them. The scene repeats itself with new tourists checking on the popularity of each restaurant before entering one

of them. After a while, all newcomers choose the same restaurant: they choose the more popular one irrespective of their own information. This tale illustrates how rational people may herd and choose one action because it is chosen by others. Among the many similar stories, two are particularly enlightening.

High sales promote high sales

In 1995, management gurus Michael Reacy and Fred Wiersema secretly purchased 50,000 copies of their business strategy book *The Discipline of Market Leaders* from stores which were monitored for the bestseller list of the *New York Times*¹. The authors must have been motivated by the following argument: people observe the sales, but not the payoffs of the purchases (assuming they have few opportunities to meet other readers). Of course, if the manipulation had been known it would have had no effect, but people rationally expect that for any given book, the probability of manipulation is small, and that the high sales must be driven by some informed buyers.

The previous story illustrates one possible motivation for using the herding effect but it is only indicative. For an actual measurement, we turn to Hanson and Putler (1996) who conducted a nice experiment which combines the control of a laboratory with a “real situation”. They manipulated a service provided by America Online (AOL) in the summer of 1995. Customers of the service could download games from a bulletin board. The games were free, but the download entailed some cost linked to the time spent in trying out the game. Some games were downloaded more than others.

The service of AOL is summarized by the window available to subscribers which is reproduced in

??: column 1 shows the first date the product was available; column 2 the name of the product, which is informative; column 4 the most recent date the file was downloaded. Column 3 is the most important and shows the number of customers who have downloaded the file so far. It presents an index of the “popularity” of the product. The main goal of the study is to investigate whether a high popularity increases the demand *ceteris paribus*.

The impact of a treatment is measured by the increase in the number of downloads per day, after the treatment, as a fraction of the average daily download (for the same product) before the treatment. The results are reported in Figure ???. All treatments have an impact and the impact of the heavy treatment (100 percent) is particularly remarkable.

¹See Bikhchandani, Hirshleifer and Welch (1998), and Business Week, August 7, 1995. Additional examples are given in Bikhchandani, Hirshleifer and Welch, (1992).




Figure 4.1: Applications for downloads

The experiment has an obvious implication for the general manipulation of demand through advertisements.

To ensure *ceteris paribus*, Hanson and Putler selected pairs of similar files which were offered by AOL. Similarity was measured by characteristics and “popularity” at a specific date. Once a pair was selected, one of the files was kept as the “control”, the other was the “treatment”. The authors boosted the popularity index of the treatment file by downloading it repeatedly. The popularity indexed was thus increased in a short session by percentage increments of 25, 50 and 100. Customers of the service were not aware that they were manipulated.

The essential issue and the framework of analysis

The previous examples share a common feature which is essential: individuals observe the actions of others and the space of actions is discrete. The actions are the words for the communication of information between agents. In the previous chapter, agents chose an action in a rich set made of all the real numbers. Here the finite number of actions exerts

a strong restriction on the vocabulary of social communication.

If there is a seminal study on social learning, it is the paper by Bikchandani, Hirshleifer and Welch (1992), hereafter BHW². They introduced the definition of informational cascades in models of Bayesian learning. In a cascade, the public belief, which is gathered from the history of observations, dominates the private signal of any individual: the action of any agent does not depend on his private information. In a cascade, all agents are herding. Since actions do not convey private informations, nothing is learned and the cascade goes on forever, possibly with an incorrect action. The failure of social learning is spectacular.

There is an essential
difference between
a cascade and a herd.

**

A cascade generates a herd, but the concepts of cascade and herd are distinct. A herd is defined as an outcome where all agents take the same action after some period. Which period is a random event and it is unknown at the start of the learning process. Actually, we we will see that it may never be known. In a herd not all agents may be herding. It is precisely because not all agents are herding in a herd that some learning takes place. The probability that the herd could be broken generates some information. But this probability must be vanishingly small for the herd to be sustained. Hence, the amount of social learning in a herd is very small.

Cascades do not occur, except in very special models which are not generic, while herds always take place eventually. The reader may think that cascades are therefore not important. Wrong: cascades are good approximations for the properties of the generic models of learning from others' actions when these actions are discrete.

**

Suppose that the set of states is finite. The *support* of a distribution of probabilities (beliefs) is the set of states that have a strictly positive belief (probability). If the set of the states is a continuum, for example a continuum of real numbers and the distribution of probabilities has a continuous distribution, the support of the the distribution is the subset where the density function is strictly positive. We will see that the process of social learning depends on whether the probability distributions, of the state or or a signal on the state, are bounded or not.

The general model is built on the models with bounded private beliefs which have been presented in Section ???. (The reader is advised to review that section if necessary). The evolution of the beliefs is presented in a diagram which will be used later in the book.

When the support is bounded, private beliefs become dominated by a public belief which is

²Banerjee (1992) presented at the same time another paper on herding, but its structure is more idiosyncratic and one cannot analyze the robustness of its properties.

either optimistic or pessimistic, as the number of observations increases. Such a situation actually never occurs when private beliefs have a distribution without points of positive mass (which is not just a perturbation of a distribution with such points). However, the limit behavior of the model is closely approximated by cascades.

Beliefs converge to the truth, almost surely, only if the support of the distribution of beliefs is unbounded. In this respect, the results of BHW have been criticized as not robust. Such theoretical focus on the limit beliefs is misleading. What matters is the speed of convergence.

Section 4 presents a detailed analysis of herds and the convergence of beliefs³. Herds always take place eventually, as a consequence of the Martingale Convergence Theorem. There is in general some learning in a herd, but that learning is very slow. The conclusions of the simple model of BHW are shown to be extraordinarily robust. They reinforce the central message of the models of learning from others which is the self-defeating property of social learning when individuals use rationally the public information.

The social optimum

In an equilibrium, no agent takes into account the externality created by his action for the information of others. In a social optimum, this externality is taken into account (as in the model with actions in a continuum, Section 3.4). A social optimum is constrained in the sense that each agent “speaks” to others only through his action. An agent has a decision rule according to which his action depends on his private belief and the public belief. He can reveal his private belief only through his action. He departs from the selfish rule of using history for his own payoff only if the externality provided to others outweighs the personal loss.

In Section ??, it is shown that the social optimal rule is to forget history if the belief from history—the public belief—is in some interval of values, and to herd otherwise. If the belief is outside of that “interval of experimentation”, there is no social learning anymore. The socially optimal rule may be implemented by setting a price of investment contingent on the public belief.

Monopoly pricing of a new good

A monopoly who captures some consumer surplus will take into account the benefit of experimentation for the future. This problem is considered in Section ??. A monopoly

³For this section, I have greatly benefited from the insights of Lones Smith and I am very grateful to him.

introduces on the market a new good of imperfectly known quality. The optimal strategy is divided in two phases. The first is the “elitist phase”: the price of the good is relatively high. Only the agents with a good signal on the good buy and the volume of sales raises the estimate of the other agents. When this estimate is sufficiently high, the monopoly lowers the price to reach all customers.

The incentive to learn is inversely related to the discount rate. If the discount rate is vanishingly small, the difference between the level of social welfare and the monopoly profit converges to zero. At the limit, the monopoly follows a strategy which is socially optimal. (Monopoly profits are redistributed).

4.1 The basic model of herding

A textbook case on
how to build a model

Students sometimes wonder how to build a model. Bikhchandani, Hirshleifer and Welsh (1992), hereafter BHW, provide an excellent lesson of methodology: (i) a good story simplifies the complex reality and keeps the main elements; (ii) this story is translated into a set of assumptions about the structure of a model (information of agents, payoff functions); (iii) the equilibrium behavior of rational agents is analyzed; (iv) the robustness of the model is examined through extensions of the initial assumptions.

We begin here with the tale of two restaurants, or a similar story where agents have to decide whether to make a fixed size investment. We construct a model with two states (defining which restaurant is better), two signal values (which generate different beliefs), and two possible actions (eating at one of two restaurants)⁴.

4.1.1 The 2 by 2 by 2 model

1. The state of nature θ has two possible values, $\theta \in \Theta = \{0, 1\}$, and is set randomly once and for all at the beginning of the first period⁵ with a probability μ_1 for the “good” state $\theta = 1$.

2. N or a countable number of agents are indexed by the integer t . Each agent’s private

⁴The example of the restaurants at the beginning of this chapter is found in Banerjee (1992). The model in this section is constructed on this story. It is somewhat mistifying that Banerjee after introducing herding through this example, develops an unrelated model which is somewhat idiosyncratic. A simplified version is presented in Exercise ??.

⁵The value of θ does not change because we want to analyze the changes in beliefs which are caused only by endogenous behavior. Changes of θ can be analyzed in a separate study (see the bibliographical notes).

information takes the form of a SBS (symmetric binary signal) with precision $q > 1/2$: $P(s_t = \theta \mid \theta) = q$.

3. Agents take an action in an *exogenous order* as in the previous models of social learning. The notation can be chosen such that agent t can make a decision in period t and in period t only. An agent chooses his action x in the discrete set $X = \{0, 1\}$. The action $x = 1$ may represent entering a restaurant, hiring an employee, or in general making an investment of a fixed size. The yield of the action x depends on the state of nature and is defined by

$$u(x, \theta) = \begin{cases} 0, & \text{if } x = 0, \\ \theta - c, & \text{if } x = 1, \text{ with } 0 < c < 1. \end{cases}$$

Since $x = 0$ or 1 , another representation of the payoff is $u(x, \theta) = (\theta - c)x$. The cost of the investment c is fixed.⁶ The yield of the investment is positive in the good state and negative in the bad state. Under uncertainty, the payoff of the agent is the expected value of $u(x, \theta)$ conditional on the information of the agent. By convention, if the payoff of $x = 1$ is zero, the agent chooses $x = 0$.

4. As in the previous models of social learning, the information of agent t is his private signal and the *history* $h_t = (x_1, \dots, x_{t-1})$ of the actions of the agents who precede him in the exogenous sequence. The *public belief* at the beginning of period t is the probability of the good state conditional on the history h_t which is public information. It is denoted by μ_t :

$$\mu_t = P(\theta = 1 \mid h_t).$$

Without loss of generality, μ_1 is the same as nature's probability of choosing $\theta = 1$.

4.1.2 Informational cascades

Agents with a good signal $s = 1$ will be called optimists and agents with a bad signal $s = 0$ will be called pessimists. An agent combines the public belief with his private signal to form his belief. If μ is the public belief in some arbitrary period, the belief of an optimist is higher than μ and the belief of a pessimist is lower. Let μ^+ and μ^- be the beliefs of the optimists and the pessimists⁷: $\mu^- < \mu < \mu^+$.

A pessimist invests if and only if his belief μ^- is greater than the cost c , *i.e.* if the public belief is greater than some value $\mu^{**} > c$. (If $c = 1/2$, $\mu^{**} = q$). If the public belief is

⁶In the tale of two restaurants, c is the opportunity cost of not eating at the other restaurant.

⁷By Bayes' rule,

$$\mu^- = \frac{\mu(1-q)}{\mu(1-q) + (1-\mu)q} < \mu < \frac{\mu q}{\mu q + (1-\mu)(1-q)} = \mu^+.$$

such that a pessimist invests, then *a fortiori*, it induces an optimist to invest. Therefore, if $\mu_t > \mu^{**}$ agent t invests whatever his signal. If $\mu_t \leq \mu^{**}$, he does not invest if his signal is bad.

Likewise, let μ^* be the value of the public belief such that $\mu^+ = c$. If $\mu_t \leq \mu^*$, agent t does not invest no matter the value of his private signal. If $\mu_t > \mu^*$ he invests if he has a good signal. The cases are summarized in the next result.

PROPOSITION 4.1. *In any period t , given the public belief μ_t :*

if $\mu^ < \mu_t \leq \mu^{**}$, agent t invests if and only if his signal is good ($s_t = 1$);*

*if $\mu_t > \mu^{**}$, agent t invests independently of his signal;*

if $\mu_t \leq \mu^$, agent t does not invest independently of his signal.*

Cascades and herds

Proposition 4.1 shows that if the public belief, μ_t , is above μ^{**} , agent t invests and ignores his private signal. His action conveys no information on this signal. Likewise, if the public belief is smaller than μ^* , then the agent does not invest. This important situation deserves a definition.

DEFINITION 4.1. *An agent herds on the public belief when his action is independent of his private signal.*

The herding of an agent describes a decision process. The agent takes into account only the public belief; his private signal is too weak to matter. If all agents herd, no private information is revealed. The public belief is unchanged at the beginning of the next period and the situation is identical: the agent acts according to the public belief whatever his private signal. The behavior of each agent is repeated period after period. This situation has been described by BHW as an *informational cascade*. The metaphor was used first by Tarde at the end of the nineteenth century.

DEFINITION 4.2. *If all agents herd (Definition 4.1), there is an informational cascade.*

We now have to make an important distinction between the herding of all agents in an informational cascade and the definition of a herd.

DEFINITION 4.3. *A herd takes place at date T if all actions after date T are identical: for all $t > T$, $x_t = x_T$.*

In a cascade, all agents are herding and make the same decision which depends only on the public belief (which stays invariant over time). Hence, all actions are identical.

PROPOSITION 4.2. *If there is an informational cascade in period t , there is a herd in the same period.*

An important distinction between an informational cascade and a herd is that a herd may occur without a cascade.

The converse of Proposition 4.2 is not true. *Herds and cascades are not equivalent.* In a herd, all agents turn out to choose the same action—in all periods—although some of them could have chosen a different action. We will see later that generically, cascades do not occur, but herds eventually begin with probability one! Why do we consider cascades then? Because their properties are stylized representations of models of social learning.

In the present model, an informational cascade takes place if $\mu_t > \mu^{**}$ or $\mu_t \leq \mu^*$. There is social learning only if $\mu^* < \mu_t \leq \mu^{**}$. Then $x_t = s_t$ and the action reveals perfectly the signal s_t . The public belief in period $t + 1$ is the same as that of agent t as long as a cascade has not started. The history of actions $h_t = (x_1, \dots, x_{t-1})$ is equivalent to the history of signals (s_1, \dots, s_{t-1}) .

Assume that there is no cascade in periods 1 and 2 and that $s_1 = 1$ and $s_2 = 1$. Suppose that agent 3 is a pessimist. Because all signals have the same precision, his bad signal “cancels” one good signal. He therefore has the same belief as agent 1 and should invest. There is a cascade in period 3.

Likewise, two consecutive bad signals ($s = 0$) start a cascade with no investment, if no cascade has started before. If the public belief μ_1 is greater than c and agent 1 has a good signal, a cascade with investment begins in period 2. If $\mu_1 < c$ and the first agent has a bad signal, he does not invest and a cascade with no investment begins in period 2.

In order *not to* have a cascade, a necessary condition is that the signals alternate consecutively between 1 and 0. We infer that

- the probability that a cascade has not started by period t converges to zero exponentially, like β^t for some parameter $\beta < 1$;
- there is a positive probability that the cascade is wrong: in the bad states all agents may invest after some period, and investment may stop after some period in the good state;

- beliefs do not change once a herd has started; rational agents do not become more confident in a cascade.

PROPOSITION 4.3. *When agents have a binary signal, an informational cascade occurs after some finite date, almost surely. The probability that the informational cascade has not started by date t converges to 0 like β^t for some β^t with $0 < \beta < 1$.*


A geometric representation

The evolution of the beliefs is represented in Figure ?? . In each period, a segment represents the distribution of beliefs: the top of the segment represents the belief of an optimist, the bottom the belief of a pessimist and the mid-point the public belief. The segments evolve randomly over time according to the observations.

In the first period, the belief of an optimist, μ_1^+ , is above c while the belief of a pessimist, μ_1^- , is below c . The action is equal to the signal of the agent and thus reveals that signal. In the figure, $s_1 = 0$, and the first agent does not invest. His information is incorporated in the public information: the public belief in the second period, μ_2 , is identical to the belief of the first agent: $\mu_2 = \mu_1^-$. The sequence of the signal endowments is indicated in the figure. *When there is social learning, the signal of agent t is integrated in the public information of period $t + 1$.* Using the notation of the previous chapter, $\mu_{t+1} = \tilde{\mu}_t$.

Consider now period 5 in the figure: agent 5 is an optimist, invests and reveals his signal since he could have been a pessimist who does not invest. His information is incorporated in the public belief of the next period and $\mu_6 = \mu_5^+$. The belief of a pessimist in period 6 is now higher than the cost c (here, it is equal to the public belief μ_5). In period 6, the belief of an agent is higher than the cost of investment, whatever his signal. He invests, nothing is learned and the public belief is the same in period 7: a cascade begins in period 6. The cascade takes place because all the beliefs are above the cut-off level c . This condition is met here because the public belief μ_6 is strictly higher than μ^{**} . Since μ_6 is identical to the belief of an optimist in period 5, the cascade occurs because the beliefs of all investing agents are strictly higher than μ^{**} in period 5. A cascade takes place because of the high belief of the last agent who triggers the cascade. Since this property is essential for the occurrence of an informational cascade, it is important and will be discussed later in more details.

In this simple model, the public belief $\mu_t = P(\theta = 1|h_t)$ converges to one of two values (depending on the cascade). From the Martingale Convergence Theorem, we knew μ_t would necessarily converge in probability. The exponential convergence is particularly fast. The informational cascade may be incorrect however: all agents may take the wrong decision.



In each period, the middle of the vertical segment is the public belief, while the top and the bottom of the segment are the beliefs of an optimist (with a private signal $s = 1$) and of a pessimist (with signal $s = 0$). The private signals are $s_1 = 0, s_2 = 1, s_3 = 0, s_4 = 1, s_5 = 1$.

Figure 4.2: Cascade representation

Black sheeps

Assume there is a cascade in some period T in which agents invest whatever their signal. Extend now the previous setting and assume that agent T may be of one of two types. Either he has a signal of precision q like the previous agents, or his precision is $q' > q$ and q' is sufficiently high with respect to the public belief that if he has a bad signal ($s_T = 0$), he does not invest. The type of the agent is private and therefore not observable, but the possibility that agent T has a higher precision is known to all agents.

Suppose that agent T does not invest: $x_T = 0$. What inference is drawn by others? The only possible explanation is that agent T has a signal of high precision q' and that his signal is bad: the information of agent T is conveyed *exactly* by his action.

If agent T invests, his action is like that of others. Does it mean that the public belief does not change? No! The absence of a black sheep in period T (who would not invest) increases the confidence that the state is good. Social learning takes place as long as not all agents herd. The learning may slow down however as agents with a relatively low precision begin to herd. The inference problem with heterogeneous precisions requires a model which incorporates the random endowment of signals with different precisions. A model with two

types of precision is presented in the appendix.

The simple model has served two useful purposes: (i) it is a lesson on how to begin to think formally about a stylized fact and the essence of a mechanism; (ii) it strengthens the intuition about the mechanism of learning and its possible failures. These steps need to be as simple as possible. But the simplicity of the model could generate the criticism that its properties are not robust. The model is now generalized and we will see that its basic properties are indeed robust.

4.2 The standard model with bounded beliefs

We now extend the previous model to admit any distribution of private beliefs as described in Section ???. Such a distribution is characterized by the *c.d.f.* $F^\theta(\mu)$ which depends on the state θ . Recall that the *c.d.f.s* satisfy the Proportional Property (??) and therefore the assumption of first order stochastic dominance: for any μ in the interior of the support of the distribution, $F^{\theta_0}(\mu) > F^{\theta_1}(\mu)$. By an abuse of notation, $F^\theta(\mu)$ will represent the *c.d.f.* of a distribution of the beliefs measured as the probability of θ_1 , and $F^\theta(\lambda)$ will represent the *c.d.f.* of a distribution of the LLR between θ_1 and θ_0 .

We keep the following structure: two states $\theta \in \{\theta_0, \theta_1\}$, two actions $x \in \{0, 1\}$, with a payoff $(E[\theta] - c)x$, $\theta_0 < c < \theta_1$. The states θ_1 and θ_0 will be called “good” and “bad”. We may take $\theta_0 = 1$ and $\theta_0 = 0$, but the notation may be clearer if we keep the symbols θ_1 and θ_0 rather than using numerical values.

4.2.1 Social learning

At the end of each period, agents observe the action x_t . Any belief λ is updated using Bayes’ rule. This rule is particularly convenient when expressed in LLR as in equation (??) which is repeated here.

$$\lambda_{t+1} = \lambda_t + \nu_t, \quad \text{with} \quad \nu_t = \text{Log}\left(\frac{P(x_t|\theta_1)}{P(x_t|\theta_0)}\right). \quad (4.1)$$

The updating term ν_t is independent of the belief λ_t . Therefore, the distribution of beliefs is translated by a random term ν_t from period t to period $t + 1$. Agent t invests if and only if his probability of the good state is greater than his cost, *i.e.* if his LLR, λ , is greater than $\gamma = \text{Log}(c/(1 - c))$. The probability that agent t invests depends on the state and is equal to $\pi_t(\theta) = 1 - F_t^\theta(\gamma)$.

Observations

	$x_t = 1$	$x_t = 0$	
States of Nature	$\theta = \theta_1$	$1 - F_t^{\theta_1}(\gamma)$	$F_t^{\theta_1}(\gamma)$
	$\theta = \theta_0$	$1 - F_t^{\theta_0}(\gamma)$	$F_t^{\theta_0}(\gamma)$

with $\gamma = \text{Log}\left(\frac{c}{1-c}\right)$.

The action in period t , $x_t \in \{0, 1\}$, provides a binary random signal on θ with probabilities described in Table ???. Since the *c.d.f.* F^{θ_1} dominates F^{θ_0} in the sense of first order stochastic dominance (Proposition ??), there are more optimistic agents in the good than in the bad state on average. Hence, the probability of investment is higher in the good state, and the observation $x_t = 1$ raises the beliefs of all agents.

Following the observation of x_t , the updating equation (??) takes the particular form

$$\lambda_{t+1} = \lambda_t + \nu_t, \quad \text{with} \quad \nu_t = \begin{cases} \text{Log}\left(\frac{1 - F_t^{\theta_1}(\gamma)}{1 - F_t^{\theta_0}(\gamma)}\right), & \text{if } x_t = 1, \\ \text{Log}\left(\frac{F_t^{\theta_1}(\gamma)}{F_t^{\theta_0}(\gamma)}\right), & \text{if } x_t = 0. \end{cases} \quad (4.2)$$

In this equation, $\nu_t \geq 0$ if $x_t = 1$ and $\nu_t \leq 0$ if $x_t = 0$. The observation of x_t conveys some information on the state as long as $F_t^{\theta_1}(\gamma) \neq F_t^{\theta_0}(\gamma)$.

Since the distribution of LLRs is invariant up to a translation, it is sufficient to keep track of one of the beliefs. If the support of beliefs is bounded, we choose the mid-point of the support, called by an abuse of notation the public belief. If the support is not bounded, the definition of the public belief will depend on the particular case.

The Markov process

The previous process has an abstract formulation which may provide some perspective on the process of social learning. We have seen that the position of the distribution in any period can be characterized by one point λ_t . Let μ_t be the belief of an agent with LLR equal to λ_t . The Bayesian formula (4.2) takes the general form $\mu_{t+1} = B(x_t, \mu_t)$ and x_t is a random variable which takes the value 1 or 0 according to Table ??. These values depend only on λ_t and therefore μ_t depends on θ . The process of social learning is summarized by

the equations

$$\begin{cases} \mu_{t+1} = B(\mu_t, x_t), \\ P(x_t = 1) = \pi(\mu_t, \theta). \end{cases} \quad (4.3)$$

The combination of the two equations defines a Markov process for μ_t . Such a definition is natural and serves two purposes. It provides a synthetic formulation of the social learning. It is essential for the analysis of convergence properties. However, such a formulation can be applied to a wide class of processes and does not highlight specific features of the structural model of social learning with discrete actions.

4.2.2 Bounded beliefs


Assume the initial distribution of private beliefs is *bounded*. Its support is restricted to a finite interval $(\underline{\lambda}_1, \bar{\lambda}_1)$. This case is represented in Figure ???. Let λ_t be the public belief in period t , *i.e.*, the mid-point of the support: $\lambda_t = (\underline{\lambda}_t + \bar{\lambda}_t)/2$ and let $\sigma = (\bar{\lambda}_t - \underline{\lambda}_t)/2$, a constant. If λ_t is greater than the value $\lambda^{**} = \gamma + \sigma$, the support of the distribution is above γ and agent t invests, whatever his belief. Likewise, if $\lambda \leq \lambda^* = \gamma - \sigma$, no agent invests. In either case, there is an informational cascade. There is no informational cascade as long as the public belief stays in the interval $(\lambda^*, \lambda^{**}) = (\gamma - \sigma, \gamma + \sigma)$. The complement of that interval will be called the *cascade set*.

Figure ??? is drawn under the assumption of an atomless distribution of beliefs but it can also be drawn with atoms as in Figure ???.

i We know from the Martingale Convergence Theorem that the probability of the good state, $\mu_t = e^{\lambda_t}/(1 + e^{\lambda_t})$, converges in probability. Hence, λ_t must converge to some value. Suppose that the limit is not in the cascade set. Then, asymptotically, the probability that $x_t = 1$ remains different in states θ_1 and θ_0 . Hence, with strictly positive probability, the common belief is updated by some non vanishing amount, thus contradicting the convergence of the martingale. This argument is used in the Appendix to prove that λ_t must converge to a value in the cascade set.

PROPOSITION 4.4. *Assume that the support of the initial distribution of private beliefs is $I = [\lambda_1 - \sigma, \lambda_1 + \sigma]$. Then λ_t converges almost surely to a limit $\lambda_\infty \notin (\gamma - \sigma, \gamma + \sigma)$ with $\gamma = \text{Log}(c/(1 - c))$.*

Is the occurrence of a cascade generic?



In each period, the support of the distribution of beliefs (LLR) is represented by a segment. The action is $x_t = 1$ if and only if the belief (LLR) of the agent is above γ . If agent t happens to have a belief above (below) γ , the distribution moves up (down) in the next period $t + 1$. If the entire support is above (below) γ , the action is equal to 1 (0) and the distribution stays constant.

Figure 4.3: Cascade representation

Generically, cascades do not occur!

The previous result shows that the beliefs tend to the cascade set. But for an arbitrary distribution of initial beliefs, is this convergence as fast as with the discrete beliefs of Figure ??, or is it slow? It turns out that, for most distributions which are “smooth”, the convergence is slow and cascades do not occur.

The mechanism can be explained simply. Suppose that the beliefs converge to the upper part of Figure ?? where agents invest. The probability that agent t invests is lower in the bad than in the good state, but as the beliefs move upwards, these two probabilities converge to each other with a common limit equal to one. The observation of an investment conveys a vanishingly small amount of information and the upward shift of the beliefs is also vanishingly small⁸.

Assume the distribution of initial beliefs has a density $f^\theta(\mu)$ in state θ such that

$$f^1(\mu) = \mu\phi(\mu), \quad \text{and} \quad f^0(\mu) = (1 - \mu)\phi(\mu), \quad (4.4)$$

⁸The argument does not apply when beliefs are high and there is no investment in the period. In that case, the probability of no investment is low in both states, but the ratio between these probabilities is not small.

for some function $\phi(\mu)$ with a support in $[a, 1 - a]$, $a > 0$. This distribution is “natural” in the sense that it is generated by a two-step process in which agents draw a SBS of precision μ with a density proportional to $\phi(\mu)$. A simple case is provided by a uniform distribution of precisions where ϕ is constant. The proof of the following result is left to the reader.

PROPOSITION 4.5. *Assume that the density of initial beliefs are proportional to μ and to $1 - \mu$ in the two states. If there is no cascade in the first period, there is no cascade in any period.*

The result applies if $\phi(\mu)$ does not put too much mass at either end of its support. This is intuitive: in the model with discrete beliefs (Figure ??), all the mass is put at either end of the support. A smooth perturbation of the discrete model does not change its properties. A numerical simulation of the case which satisfies (4.4) shows that a cascade occurs if $\phi(\mu) = x^{-n}$ with n sufficiently high ($n \geq 4$ for a wide set of other parameters). In this case, the distribution puts a high mass at the lower end of the support.

Right and wrong cascades

A cascade may arise with an incorrect action: for example, beliefs may be sufficiently low that no agent invests while the state is good. However, agents learn rationally and the probability of a wrong cascade is small if agents have a wide diversity of beliefs as measured by the length of the support of the distribution.

Suppose that the initial distribution in LLR is symmetric around 0 with a support of length 2σ . We compute the probability of a wrong cascade for an agent with initial belief $1/2$. A cascade with no investment arises if his LLR λ_t is smaller than $\gamma - \sigma$, *i.e.*, if his belief in level is such that

$$\mu_t \leq \epsilon = e^{\gamma - \sigma} / (1 + e^{\gamma - \sigma}).$$

When the support of the distribution in LLR becomes arbitrarily large, $\sigma \rightarrow \infty$ and ϵ is arbitrarily small. From Proposition ?? with $\mu_1 = 1/2$, we know that

$$P(\mu_t \leq \epsilon | \theta_1) \leq 2\epsilon.$$

The argument is the same for the cascades where all agents invest. The probability of a wrong cascade for a neutral observer (with initial belief $1/2$) tends to zero if the support of the distribution in LLR becomes arbitrarily large (or equivalently if the beliefs measured as probabilities of θ_1 are intervals converging to $(0, 1)$).

PROPOSITION 4.7. *If the support of the initial distribution of LLRs contains the interval $[-\sigma, +\sigma]$, then for an observer with initial belief $1/2$, the probability of a wrong cascade is less than 4ϵ , with $\epsilon = e^{-\sigma}c / (1 - c + e^{-\sigma}c)$.*

4.3 The convergence of beliefs

When private beliefs are bounded, beliefs never converge to perfect knowledge. If the public belief would converge to 1 for example, in finite time it would overwhelm any private belief and a cascade would start thus making the convergence of the public belief to 1 impossible. This argument does not hold if the private beliefs are unbounded because in any period the probability of a “contrarian agent” is strictly positive.

4.3.1 Unbounded beliefs: convergence to the truth

From Proposition 4.7 (with $\sigma \rightarrow \infty$), we have immediately the next result.

PROPOSITION 4.8. *Assume that the initial distribution of private beliefs is unbounded. Then the belief of any agent converges to the truth: his probability assessment of the good state converges to 1 in the good state and to 0 in the bad state.*

Does convergence to the truth matter?

A bounded distribution of beliefs is necessary for a herd on an incorrect action, as emphasized by Smith and Sørensen (1999). Some have concluded that the properties of the simple model of BHW are not very robust: cascades are not generic and do not occur for sensible distributions of beliefs; the beliefs converge to the truth if there are agents with sufficiently strong beliefs. In analyzing properties of social learning, the literature has often focused on whether learning converges to the truth or not. This focus is legitimate for theorists, but it is seriously misleading. What is the difference between a slow convergence to the truth and a fast convergence to an error? From a welfare point of view and for many people, it is not clear.

To focus on whether social learning converges to the truth or not can be misleading.


The focus on the ultimate convergence has sometimes hidden the central message of studies on social learning: the combination of history’s weight and of self-interest slows down the learning from others. The beauty of the BHW model is that it is non generic in some sense (cascades do not occur under some perturbation), but its properties are generic.

If beliefs converge to the truth, the speed of convergence is the central issue. This is why the paper of Vives (1993) has been so useful in the previous chapter. We learned from that model that an observation noise reduces the speed of the learning from others. Since the discreteness of the action space is a particularly coarse filter, the slowing down of social learning should also take place here. When private beliefs are bounded, the social learning does not converge to the truth. When private beliefs are unbounded, we should observe a slow rate of convergence.

We saw that cascades do not occur for sensible distributions of beliefs because the signal of the action (investment or no investment) is vanishingly weak when the public belief tends to the cascade set corresponding to the action. This argument applies when the distribution of beliefs is unbounded, since the mass of atoms at the extreme ends of the distribution must be vanishingly small. Hence, there is an immediate presumption that social learning must be slow asymptotically. The slow learning is first illustrated in an example and then analyzed in detail.

A numerical example

The private signals are defined by $s = \theta + \epsilon$ where ϵ is normally distributed with variance σ^2 . An exercise shows that if μ tends to 0, the mass of agents with beliefs above $1 - \mu$ tends to zero faster than any power of μ . A numerical example of the evolution of beliefs is presented in Figure ???. One observes immediately that the pattern is similar to a cascade in the BHW model with the occurrence of “black sheeps”.



The upper graph represents the evolution of the public belief. The lower graph represents the sequence of individuals' actions. It is distinct from the horizontal axis only if $x_t = 1$.

For this example only, it is assumed that the true state is 1. The initial belief of the agent is $\mu_1 = 0.2689$, (equivalent to a LLR of -1), and $\sigma = 1.5$. The actions of individuals in each period are presented by the lower schedule (equal to 0.1 if $x_t = 1$ and to 0 otherwise). For the first 135 periods, $x_t = 0$ and μ_t decreases monotonically from around 0.27 to around 0.1. In period 136, the agent has a signal which is sufficiently strong to have a belief $\tilde{\mu}_{136} > c = 0.5$ and he invests. Following this action, the public belief is higher than 0.5

(since 0.5 is a lower bound on the belief of agent 135), and $\mu_{137} > 0.5$. In the example, $\mu_{137} = 0.54$. The next two agents also invest and $\mu_{139} = 0.7$. However, agent 139 does not invest and hence the public belief must fall below 0.5: $\mu_{140} = 0.42$. Each time the sign of $\mu_{t+1} - \mu_t$ changes, there is a large jump in μ_t .

The figure provides a nice illustration of the herding properties found by BHW in a model with “black sheeps” which deviate from the herds. The figure exhibits two properties which are standard in models of social learning with discrete decisions:

- (i) when μ_t eventually converges monotonically to the true value of 1 (after period 300 here), the convergence is very slow;
- (ii) when a herd stops, the public belief changes by a quantum jump.

The slow learning from others

Assume now a precision of the private signals such that $\sigma_\epsilon = 4$, and an initial public belief $\mu_1 = 0.2689$ (with a LLR equal to -1). The true state is good. The model was simulated for 500 periods and the public belief was computed for period 500. The simulation was repeated 100 times. In 97 of the 100 simulations, no investment took place and the public belief decreased by a small amount to a value $\mu_{500} = 0.2659$. In only three cases did some investment take place with μ_{500} equal to 0.2912, 0.7052 and 0.6984, respectively. Hardly a fast convergence!

By contrast, consider the case where agents observe directly the private signals of others and do not have to make inferences from the observations of private actions. From the specification of the private signals and Bayes’ rule,

$$\lambda_{t+1} = \lambda_1 + t \left(\frac{\theta_1 - \theta_0}{\sigma_\epsilon^2} \right) \left(\frac{\theta_1 - \theta_0}{2} + \eta_t \right), \quad \text{with} \quad \eta_t = \frac{1}{t} \sum_{k=1}^t \epsilon_k.$$

Given the initial belief $\mu_1 = 0.2689$, $\theta_0 = 0$, $\theta_1 = 1$, $t = 499$ and $\sigma_\epsilon = 4$,

$$\lambda_{500} = -1 + (31.2)(0.5 + \eta_{500}),$$

where the variance of η_{500} is $16/499 \approx (0.18)^2$. Hence, λ_{500} is greater than 5.33 with probability 0.95. Converting the LLR in probabilities, μ_{500} belongs to the interval (0.995, 1) with probability 0.95. What a difference with the case where agents observed private actions! The example—which is not particularly convoluted—shows that the convergence to the truth with unbounded private precisions may not mean much practically. Even when the distribution of private signals is unbounded, the process of social learning can be very slow when agents observe discrete actions. Cascades provide a better stylized description

of the properties of social learning through discrete actions than the convergence result of Proposition 4.8. The properties of the example are confirmed by the general analysis of the convergence in Section ??.

4.4 Herds and the slow convergence of beliefs

4.4.1 Herds

The Martingale Convergence Theorem implies that the public belief converges almost surely. Assume that the distribution of beliefs is bounded. At the limit, the support of the distribution must be included in one of the two cascade sets. Suppose that on some path the support of the distribution converges to the upper half of the cascade set where all agents invest: $\underline{\mu}_t \rightarrow c$. We now prove by contradiction that the number of periods with no investment is finite on this path.

Since there is a subsequence $x_n = 0$, we may assume $\underline{\mu}_n < c$. Following the observation of $x_n = 0$, Bayes' rule implies

$$\lambda_{n+1} = \lambda_n + \nu_n, \quad \text{with} \quad \nu_n = \text{Log} \left(\frac{F^1(\lambda_1 + z_n)}{F^0(\lambda_1 + z_n)} \right), \quad \text{and} \quad z_n = \gamma - \lambda_n.$$

By the assumption of first order stochastic dominance, if $z_n \rightarrow 0$, there exists $\alpha < 0$ such that $\eta_n < \alpha$, which contradicts the convergence of λ_n : the jump down of the LLR contradicts the convergence. The same argument can be used in the case of an unbounded distribution of beliefs.

All paths with social learning end with a herd.

THEOREM 4.1. *On any path $\{x_t\}_{t \geq 1}$ with social learning, a herd begins in finite time. If the distribution of beliefs is unbounded and $\theta = \theta_1$ ($\theta = \theta_0$), there exists T such that if $t > T$, $x_t = 1$ ($x_t = 0$), almost surely.*

This result is due to Smith and Sørensen (2001). It shows that herds take place eventually although, generically, not all agents are herding in any period!

4.4.2 The asymptotic rate of convergence is zero

When beliefs are bounded, they may converge to an incorrect value with a wrong herd. The issue of convergence speed makes sense only if beliefs are unbounded. This section provides a general analysis of the convergence in the binary model. Without loss of generality, we assume that the cost of investment is $c = 1/2$.

Suppose that the true state is $\theta = 0$. The public belief μ_t converges to 0. However, as $\mu_t \rightarrow 0$, there are fewer and fewer agents with a sufficiently high belief who can go against the public belief if called upon to act. Most agents do not invest. The probability that an investing agent appears becomes vanishingly small if μ tends to 0 because the density of beliefs near 1 is vanishingly small if the state is 0. It is because no agent acts contrary to the herd, although there could be some, that the public belief tends to zero. But as the probability of contrarian agents tends to zero, the social learning slows down.

Let f^1 and f^0 be the density functions in states 1 and 0. From the proportional property (Section ??), they satisfy

$$f^1(\mu) = \mu\phi(\mu), \quad f^0(\mu) = (1 - \mu)\phi(\mu), \quad (4.5)$$

where $\phi(\mu)$ is a function. We will assume, without loss of generality, that this function is continuous.

If $\theta = 0$ and the public belief converges to 0, intuition suggests that the convergence is fastest when a herd takes place with no investment. The next result which is proven in the Appendix characterizes the convergence in this case.

PROPOSITION 4.9. *Assume the distributions of private beliefs in the two states satisfy (4.5) with $\phi(0) > 0$, and that $\theta = 0$. Then, in a herd with $x_t = 0$, if $t \rightarrow \infty$, the public belief μ_t satisfies asymptotically the relation*

$$\frac{\mu_{t+1} - \mu_t}{\mu_t} \approx -\phi(0)\mu_t,$$

and μ_t converges to 0 like $1/t$: there exists $\alpha > 0$ such that if $\mu_t < \alpha$, then $t\mu_t \rightarrow a$ for some $a > 0$.

If $\phi(1) > 0$, the same property applies to herds with investment, *mutatis mutandis*.

The previous result shows that in a herd, the asymptotic rate of convergence is equal to 0.

The domain in which $\phi(\mu) > 0$ represents the support of the distribution of private beliefs. Recall that the convergence of social learning is driven by the agents with extreme beliefs. It is therefore important to consider the case where the densities of these agents are not too small. This property is embodied in the inequalities $\phi(0) > 0$ and $\phi(1) > 0$. They represent a property of a *fat tail* of the distribution of private beliefs. If $\phi(0) = \phi(1)$, we will say that the distributions of private beliefs have *thin tails*. The previous proposition assumes the case of fat tails which is the most favorable for a fast convergence.

We know from Theorem 4.1 that a herd eventually begins with probability 1. Proposition 4.9 characterized the rate of convergence in a herd and it can be used to prove the following result¹⁰.

THEOREM 4.2. *Assume the distributions of private beliefs satisfy (4.5) with $\phi(0) > 0$ and $\phi(1) > 0$. Then μ_t converges to the true value $\theta \in \{0, 1\}$ like $1/t$.*

The benchmark: learning with observable private beliefs

When agents observe beliefs through actions, there is a loss of information which can be compared with the case where private beliefs are directly observable. In Section ??, the rate of convergence is shown to be exponential when agents have binary private signals. We assume here the private belief of agent t is publicly observable. The property of exponential convergence is generalized by the following result.

PROPOSITION 4.10. *If the belief of any agent t is observable, there exists $\gamma > 0$ such that $\mu_t = e^{-\gamma t} z_t$ where z_t tends to 0 almost surely.*

The contrast between Theorem 4.2 and Proposition 4.10 shows that the social learning through the observation of discrete actions is much slower, “exponentially slower¹¹”, than if private informations were publicly observable.

4.4.3 Why do herds occur?

Herds must eventually occur as shown in Theorem 4.1. The proof of that result rests on the Martingale Convergence Theorem: the break of a herd induces a large change of the beliefs which contradicts the convergence. Lones Smith has insisted, quite rightly, that one should provide a direct proof that herds take place for sure eventually. This is done by computing the probability that a herd is broken in some period after time t . Such a probability tends to zero as shown in the next result.

THEOREM 4.3. *Assume the distributions of private beliefs satisfy (4.5) with $\phi(0) > 0$ and $\phi(1) > 0$. Then the probability that a herd has not started by date t tends to 0 like $1/t$.*

¹⁰See Chamley (2002).

¹¹Smith and Sørensen (2001) provide a technical result (Theorem 4) which states that the Markov process defined in (4.3) exhibits exponential convergence of beliefs to the truth under some differentiability condition. Since the result is in a central position in a paper on social learning, and they provide no discussion about the issue, the reader who is not very careful may believe that the convergence of beliefs is exponential in models of social learning. Such a conclusion is the very opposite of the central conclusion of all models of learning from others’ actions. The ambiguity of their paper on this core issue is remarkable. Intuition shows that beliefs cannot converge exponentially to the truth in models of social learning. In all these models, the differentiability condition of their Theorem 4 is not satisfied (Exercise ??).

4.4.4 Discrete actions and the slow convergence of beliefs

The assumption of a “fat tail” of the distribution of beliefs, $\phi(0) > 0, \phi(1) > 0$, is easy to draw mathematically but it is not supported by any strong empirical evidence.

The thinner the tail of the distribution of private beliefs, the slower the convergence of social learning. However, if private signals are observable, the convergence is exponential for any distribution. The case of a thin tail provides a transition between a distribution with a thick tail and a bounded distribution where the convergence stops completely in finite time, almost surely (Chamley, 2002).

It is reasonable to consider the case where the density of beliefs is vanishingly small when the belief approaches perfect knowledge. We make the following assumption. For some $b > 0, c > 0$,

$$f^1(1) = 0, \quad \text{and} \quad \lim_{\mu \rightarrow 0} \left(f^1(\mu) / (1 - \mu)^b \right) = c > 0. \quad (4.6)$$

The higher is b , the thinner is the tail of the distribution near the truth. One can show that the sequence of beliefs with the history of no investment tends to 0 like $1/t^{1/(1+b)}$ (Exercise ??).

The main assumption in this chapter is, as emphasized in BHW, that actions are discrete. To simplify, we have assumed two actions, but the results could be generalized to a finite set of actions. The discreteness of the set of actions imposes a filter which blurs more the information conveyed by actions than the noise of the previous chapter where agents could choose action in a continuum. Therefore, the reduction in social learning is much more significant in the present chapter than in the previous one.

Recall that when private signals can be observed, the convergence of the public belief is exponential like $e^{-\alpha t}$ for some $\alpha > 0$. When agents choose an action in a continuum and a noise blurs the observation, as in the previous chapter, the convergence is reduced to a process like $e^{-\alpha t^{1/3}}$. When actions are discrete, the convergence is reduced, at best, to a much slower process like $1/t$. If the private signals are Gaussian, (as in the previous chapter), the convergence is significantly slower as shown in the example of Figure ??. The fundamental insight of BHW is robust.

4.6 Crashes and booms

The stylized pattern of a herd which is broken by a sudden event is emblematic of a pattern of “business as usual” where at first beliefs change little, then some event generates a crash or a boom, after which the new beliefs seem “obvious” in a “wisdom after the facts”. This

sequence has been illustrated by Caplin and Leahy (1994). Their assumption of endogenous timing is not necessary for the property.

In each period, there is a new population of agents which forms a continuum of mass one. Each agent has a private information on θ in the form of a Gaussian signal $s_t = \theta + \epsilon_t$, where ϵ_t has a normal distribution $\mathcal{N}(0, \sigma_\epsilon^2)$ and is independent of other variables. Each agent chooses a zero-one action $x \in \{0, 1\}$.

In period t , agents know the history $h_t = \{Y_1, \dots, Y_{t-1}\}$ of the aggregate variable $Y_t = X_t + \eta_t$, where X_t is the mass of investments by the agents in period t and η_t is a noise which is distributed $\mathcal{N}(0, \sigma_\eta^2)$.

If λ_t is the public LLR between states θ_1 and θ_0 , an agent with private signal s has a LLR equal to

$$\lambda(s) = \lambda_t + \frac{\theta_1 - \theta_0}{\sigma_\epsilon^2} \left(s - \frac{\theta_0 + \theta_1}{2} \right).$$

Given the net payoffs in the two states, the agent invests if and only if he believes state θ_1 to be more likely than state θ_0 , hence if his LLR is positive. This is equivalent to a private signal s such that

$$s > s^*(\lambda_t) = \frac{\theta_0 + \theta_1}{2} - \frac{\sigma_\epsilon^2}{\theta_1 - \theta_0} \lambda_t.$$

Let $F(\cdot; \sigma)$ be the *c.d.f.* of the Gaussian distribution $\mathcal{N}(0, \sigma^2)$. Since the mass of the population in period t is one, the level of aggregate endogenous investment is

$$X_t = 1 - F(s^*(\lambda_t) - \theta; \sigma_\epsilon).$$

The level of aggregate activity

$$Y_t = 1 - F(s^*(\lambda_t) - \theta; \sigma_\epsilon) + \eta_t.$$

is a noisy signal on θ . The derivative of Y_t with respect to θ is

$$\frac{\partial Y_t}{\partial \theta} = (s^*(\lambda_t) - \theta) \exp\left(-\frac{(s^*(\lambda_t) - \theta)^2}{2\sigma_\epsilon^2}\right).$$


If the cut-off point $s^*(\lambda_t)$ is far to the right or to the left, the multiplier of θ on Y_t is small and the impact of θ on Y_t is dwarfed by the observation noise η_t , exactly as in the model of Vives (1993). Hence, the information content of the observation Y_t is small when most agents invest ($s^*(\lambda_t)$ is low), or most do not invest ($s^*(\lambda_t)$ is high).

Crash and boom

Suppose that the true state is θ_0 and that the level of optimism, as measured by the LLR, is high. Most agents invest and the aggregate activity is dominated by the noise.

However, the beliefs of agents are unbounded and the public belief converges to the true state. When the public belief decreases to the “middle range”, the difference between the mass of agents in the two states becomes larger and this difference dominates the noise. The level of aggregate activity is more informative. Since the true state is θ_0 , the public belief decreases rapidly and the aggregate activity falls drastically. A crash occurs.

This property is illustrated by a simulation. Two realizations of the observation shocks are considered. In the first, all realizations of the shocks are set at zero, $\eta_t \equiv 0$. The evolution of the public belief¹⁵, measured in LLR, is represented in Part A of Figure 4. In Part B, the evolution of the public belief is represented for random realizations η_t .



Parameters: $\theta_0 = 0, \theta_1 = 1, \sigma_\eta = 0.5, \sigma_\epsilon = 2, c = 0.5, \theta = \theta_1, \mu_1 = 0.75$.

The public belief is measured in LLR on the vertical axis. (Note that the belief is not measured by the probability of state θ_1). The period is reported on the horizontal axis.

Figure 4.4: Examples of evolution of beliefs

In Part A, the public belief evolves slowly at first, then changes rapidly in a few periods and evolves slowly after. The LLR tends to $-\infty$, but this convergence is obviously very slow.

If agents could observe directly the signals of others, the curve in Part A would be replaced by a straight line. The sudden change occurs here because of the non-linearity of the

¹⁵The update of the public belief from λ_t to λ_{t+1} is given by Bayes' rule:

$$\lambda_{t+1} = \lambda_t + \text{Log} \left(\frac{f(x_t - (1 - F(s^*(\mu_t) - \theta_1; \sigma_\epsilon)); \sigma_\eta)}{f(x_t - (1 - F(s^*(\mu_t) - \theta_0; \sigma_\epsilon)); \sigma_\eta)} \right),$$

where f is the density function associated to the *c.d.f.* F .

information content of individual actions.

In Part B, the changes λ_t are also sudden. Note that even if the public belief is pessimistic (with $\lambda_t < 0$ at around $t = 100$), a string of positive shocks can induce a sudden upward jump of the public belief.

The model generates symmetrically crashes and booms. If the initial level of pessimism is low and the true state is high, eventually agents learn about it and the learning process goes through a phase of rapid changes of beliefs.

4.7 Bibliographical notes

Social learning in a changing world

Throughout this chapter and the next, the state of nature is invariant. This assumption is made to focus on the learning of a given state and it applies when the state does not change much during the phase of learning. Assume now, following Moscarini, Ottaviani and Smith (1998) that the value of θ switches between θ_0 and θ_1 according to a random Markov process: the set of states of nature $\Theta = \{\theta_0, \theta_1\}$ is fixed but between periods, θ switches to the other value with probability π .

Suppose that all agents are herding in period t . Does the public belief stay constant as in the previous sections of this chapter? Agents learn nothing from the observation of others, but they know that θ evolves randomly. Ignoring the actions from others, the public belief (probability of state θ_1) regresses to the mean, $1/2$. Therefore, after a finite number of periods, the public belief does not dominate the belief of some agents in which case not all agents herd. The herding by all agents stops. This property is interesting only if π is neither too small nor too high: if π is very small, the regression to the mean is slow and the herding behavior may last a long time; if π is sufficiently large, the expectation of the exogenous change between periods is so large that the learning from others' actions which is driven by their information about past values of θ bears no relation with the current value of θ . No cascade can occur.

Experiments

The BHW model has been experimented in the laboratory by Anderson and Holt (1996), (1997). Such experiments raise the issues of the actual understanding of Bayesian inference by people (Holt and Anderson, 1996), and of the power of the tests. A important difficulty is to separate the rational Bayesian learning from *ad hoc* rules of decision making after the

observations of others' actions (such as counting the number of actions of a given type in history, or taking into account the last observed action)¹⁶. Huck and Oechssler (1998) find that the tests of Anderson and Holt are not powerful against simple rules. More recent experimental studies include Çelen and Kariv (2002b), (2002c), or Holt (2001).

¹⁶This issue is raised again in empirical studies on the diffusion of innovations (Section ??).

EXERCISE 4.1. (Probability of a wrong cascade)

Consider the $2 \times 2 \times 2$ model that we have seen in class (2 states 1 and 0, 2 actions and symmetric binary signal), where μ_1 is the prior probability of the state 1, $c \in (0, 1)$ the cost of investment, and q the precision of the binary signal. There is a large number of agents who make a decision in a fixed sequence and who observe the actions of past agents. Assume that $\mu_1 < c$ and that the difference $c - \mu_1$ is small. Let $x_t \in \{0, 1\}$ the action of agent t . We assume that the true state (unknown by agents) is $\theta = 0$.

1. Represent on a diagram with time (horizontal axis) and the probability of state 1 in the public information (vertical axis), different examples of trajectories of the public belief that end in a cascade with investment, which is a “wrong” cascade (since the state is 0). We want to compute the probability of all these wrong cascades.
2. Assume that $\theta = 1$. What is a wrong cascade?
3. Suppose that $x_1 = 0$. At the end of period 1 and beginning in period 2, what is the probability to have a wrong cascade?
4. Suppose that $x_1 = x_2 = 1$. What is the probability to have eventually a wrong cascade (with action 0)?
5. Let π_0 and π_1 be the probabilities to have a wrong cascade before any observation (at the beginning of period 1) and after the observation $x_1 = 1$. Remember that all this is conditional on $\theta = 1$. Using the geometric figure in the first question and your previous answers, show that after the observations $x_1 = 1$ and $x_2 = 0$, the probability of to end with a wrong cascade is equal to π_0 .
6. Using your previous answers, find two linear equations between π_0 and π_1 , and solve for π_0 .
7. Comment on the relation between π_0 and μ .

EXERCISE 4.2. (The model of Banerjee, 1992)

Assume that the state of nature is a real number θ in the interval $(0, 1)$, with a uniform distribution. There is a countable set of agents, with private signals equal to θ with probability $\beta > 0$, and to a number uniformly distributed on the interval $(0, 1)$ with probability $1 - \beta > 0$. (In this case the signal is not informative). The agent observes only the value of his private signal. Each agent t chooses in period t an action $x_t \in (0, 1)$. The payoff is 1 if $x_t = \theta$, and 0 if $x_t \neq \theta$. Agent t observes the history of past actions and maximizes his expected payoff. If there is more than one action which maximizes his expected payoff, he chooses one of these actions with equal probability.

1. Analyze how herds occur in this model.
2. Can a herd arise on a wrong decision?

EXERCISE 4.3. (Action set is bounded below, Chari and Kehoe, 2000)

In the standard model of this chapter, assume that agent t chooses an investment level x_t which can be any real positive number. All agents have a binary private signal with precision $p > 1/2$ and a payoff function

$$u(x, \theta) = 2(\theta - c)x - x^2, \quad \text{with } x \geq 0.$$

1. Can an informational cascade take place with positive investment? Can there be an informational cascade with no investment?
2. Show that there is a strictly positive probability of under-investment.